

Using history matching for prior choice

Wang Xueou¹, David J. Nott^{*1}, Christopher C. Drovandi², Kerrie Mengersen^{2,3}, and Michael Evans⁴

¹Department of Statistics and Applied Probability, National University of Singapore, Singapore 117546

²School of Mathematical Sciences, Queensland University of Technology, Australia

³ACEMS, the ARC Centre of Excellence in Mathematical and Statistical Frontiers

⁴Department of Statistics, University of Toronto, Toronto, Ontario, M5S 3G3, Canada

Abstract

It can be important in Bayesian analyses of complex models to construct informative prior distributions which reflect knowledge external to the data at hand. Nevertheless, how much prior information an analyst is able to elicit from an expert for use in prior construction will be limited for practical reasons, with checks for model adequacy and prior-data conflict an essential part of the process of model building and sensitivity analysis. This paper develops effective numerical methods for exploring reasonable choices of a prior distribution from a parametric class, when prior information is specified in the form of some limited constraints on prior predictive distributions, and where these prior predictive distributions are analytically intractable. The methods developed may be thought of as a novel application of the ideas of history matching, a technique developed in the literature on assessment of computer models. We illustrate the approach in the context of logistic regression and sparse signal shrinkage prior distributions for high-dimensional linear models.

^{*}Corresponding author: standj@nus.edu.sg

Keywords: Approximate Bayesian computation, Bayesian inference, History matching, Prior elicitation.

1 Introduction

Elicitation of a prior distribution is an important part of Bayesian analysis. However, often a detailed representation of an expert's beliefs is difficult to obtain, assuming it is reasonable to suppose that there are true probabilities representing an expert's beliefs at all. Even if it were possible to perform comprehensive elicitations in complex multivariate situations, it might not be worth the cost involved in many cases. In complex models, how much prior information can be easily elicited from an expert for use in prior construction will necessarily be limited for practical reasons, with checking for prior-data conflict at the analysis stage an important part of the process of model building and sensitivity analysis. For an overview of modern prior elicitation methods including realistic goals of the process, ways of evaluating its success, and the cognitive biases that make it difficult see Garthwaite et al. (2005), O'Hagan et al. (2006), Daneshkhah and Oakley (2010), Martin et al. (2012), Simpson et al. (2015) and Morris et al. (2014), among others. For a recent discussion of model checking including criticism of the prior see Chapter 5 of Evans (2015).

Here we consider the problem of predictive elicitation, where prior information is given in terms of certain limited constraints on prior predictive distributions which are not analytically tractable. By limited constraints we mean that the given prior information might rule out some distributions as unsuitable for the prior, but the prior information does not identify a unique suitable prior distribution. We will be concerned with developing effective numerical methods for identifying a reasonable value or set of values for a prior hyperparameter so that the prior satisfies the constraints. It is not our intention in this manuscript to consider the best ways to elicit the predictive constraints from an expert - these are assumed to be

given - and the numerical methods discussed here are a tool to be used as part of an iterative process of questioning and feedback in the elicitation context. A more comprehensive discussion of elicitation methods is given in the references above.

The method we propose can be thought of as a novel application of the method of history matching (Craig et al. 1997) used in the literature on assessment of computer models. A recent application of history matching in the context of a complex infectious diseases model that describes the history matching approach is Andrianakis et al. (2015). We delay further discussion of the relevant literature to Section 3. Computer models, sometimes called “simulators”, are complex computer codes that take certain inputs or parameters and produce an output. The models can either be stochastic or deterministic. The goal of history matching is to eliminate regions of the computer model parameter space where predictions from the computer model are clearly inconsistent with observed data. This may result in the conclusion that there are no plausible values of the parameters given the level of model discrepancy considered to be reasonable, and the results of a history match can guide model development and make any subsequent calibration of the model more efficient.

To apply history matching to the problem of prior choice, we can identify the computer model parameters with the prior hyperparameters to be chosen, and the computer model outputs with certain characteristics of the prior predictive densities. From these outputs an implausibility measure of the type used in history matching can be constructed. Similar to the computer models context, the approach can give an indication that there are no priors within the class considered satisfying the stated predictive constraints, as well as exploring the set of possible prior choices when the set of constraints allow for a number of suitable priors. The set of appropriate prior choices returned by the method can be used as a basis for making a unique prior choice less arbitrary, as a starting point for adding further information, or in a sensitivity analysis.

The method we discuss here, while focusing on computational problems, is in the tradi-

tion of predictive elicitation methods which elicit information about potentially observable data, rather than eliciting information about parameters directly. Examples of predictive elicitation methods in the literature for particular models include, for example, Kadane et al. (1980) and Garthwaite and Dickey (1988) for linear models, and Bedrick et al. (1996) for generalized linear models, among many others. Another popular method for informative prior choice in this tradition is the “power prior” approach of Ibrahim and Chen (2000), where a tempered version of the likelihood for data from a past study is used as the basis for the prior; if no past study is available the data can also be imaginary data created by an expert. Extensions or modifications of the method include Neuenschwander et al. (2009) and the commensurate priors of Hobbs et al. (2011). However, as mentioned above, we do not focus here on best ways to elicit prior information for particular models, either predictively or on the parameters directly. Rather, we are concerned with algorithms for finding good priors satisfying stated prior predictive constraints already given and where the relevant prior predictive distributions are analytically intractable.

A simple expository example illustrates the main features of our approach. Suppose we are to observe a binomial random variable $y \sim \text{Binomial}(n, p)$ and we are interested in inference about $p \in (0, 1)$. We parametrize the model in terms of $\beta = \log p/(1-p)$ and decide to choose a normal family for the prior, $N(0, \sigma_\beta^2)$, where σ_β^2 is to be chosen. We can think of the binomial model with this parametrization as a logistic regression with only an intercept. A less trivial logistic regression example is developed in Section 5.1. Naïvely it might be expected that setting σ_β^2 large would result in a non-informative prior. However, this is not the case as this would put most of the prior mass on large values of $|\beta|$ which correspond to values of p near 0 and 1. Setting σ_β^2 small, on the other hand, results in most of the prior mass for β near 0, which corresponds to $p = 0.5$. So both a large value of σ_β^2 , as well as a small value, would usually not be suitable for this problem - the choice of σ_β^2 requires some care. It is also clear that when n is small, so that there is little information in the data,

combining what is learnt from the data with prior information may be very important, so that a non-informative prior choice would not be desirable. Our logistic regression example in Section 5.1 illustrates the difference that even some limited prior information can make to inference in a real example.

There are a variety of ways that prior information is formulated in the elicitation literature. In our expository example and in view of the observation that a too diffuse prior would lead to the prior for p concentrating on 0 or 1, we might consider the following requirement for the prior. First, let $\hat{p} = y/n$ be the maximum likelihood estimator of p , and define the summary statistic $S = S(y) = \hat{p}(1 - \hat{p})/n$, which is an estimate of the variance of \hat{p} . If p is close to 0 or 1, we would expect \hat{p} to be close to 0 or 1 and S to be small, so if the prior predictive for S concentrates on 0, this indicates the prior is putting most of its mass near near values for p of 0 or 1. For some suitably chosen small value of S , we might require that this value be implausible under the prior predictive distribution for S and so rule out such a prior. In this simple example it might be more natural to specify prior information on the parameter p directly, but in more complex examples prior information may be more easily expressed predictively in terms of observables as we have done here. The information we have specified in this case falls short of completely determining a prior, but the methods of this paper give ways of exploring prior hyperparameter choices compatible with such information that is easily specified and thought to be important. If the analyst feels that the accuracy of any specified prior information is questionable, then, as in any Bayesian analysis, the prior should be checked to see if it conflicts with the likelihood as a part of assessing sensitivity of inferences to the prior.

In the next section we describe the basic way that we specify predictive information in the later examples and review relevant concepts of Bayesian predictive model checking since the results of certain model checks for hypothetical data summaries are the way that we formulate predictive constraints. Section 3 gives a brief introduction to the literature on history

matching and regression ABC methods. Section 4 then discusses the new approach using history matching and regression ABC for prior choice. Section 5 describes some examples and Section 6 concludes.

2 Prior information and Bayesian model checks

Consider, for a parameter of interest θ , a class of prior distributions $p(\theta|\lambda)$ indexed by a hyperparameter λ . The problem of prior choice is to choose λ . In predictive elicitation the choice will be based on some characteristics of prior predictive distributions of data or summaries of the data. Here we will describe one useful way of formulating predictive constraints for elicitation purposes, and certainly there may be others. The idea is to use the results of model checks for specified hypothetical data as a way of defining what it means for a prior elicitation to be good enough. In a sense, we treat the problem of elicitation as one of model checking (for hypothetical data).

Suppose there are some summary statistics $S^j = S^j(y)$, $j = 1, \dots, k$ of the hypothetical data y , with density $p(y|\theta)$, and that for these summary statistics we are able to say for each one whether certain values should be considered plausible or not under the prior if they were to be observed. For S^j we have a vector h^j of hypothetical values for S^j supplied by an expert, which we partition as $h^j = (h_I^j, h_P^j)$, where h_I^j is a vector of values considered as implausible by the expert, and h_P^j is a vector of values considered to be plausible. We write d_I^j for the length of h_I^j , d_P^j for the length of h_P^j , $d^j = d_I^j + d_P^j$ and $d = \sum_{j=1}^k d^j$ for the total number of constraints.

In the expository example of the introduction, we considered a $\text{Binomial}(n, p)$ model parametrized through $\beta = \log p/(1 - p)$ and $\beta \sim N(0, \sigma_\beta^2)$. Our suggested summary statistic for the elicitation was the estimated variance of the MLE, $\hat{p}(1 - \hat{p})/n$ where $\hat{p} = y/n$, and a prior predictive distribution concentrated near zero would indicate an inappropriately large

value for σ_β^2 as this corresponds to most of the prior mass on p being near 0 or 1. A suitably small implausible value for the summary here could be obtained by determining a quantile of the summary when the true p is close to 0 or 1, say 0.01 or 0.99.

We need to be precise about what plausible and implausible is. The meaning of these terms will be in terms of the result of a prior predictive check (Box 1980). Let $p(S^j|\lambda)$ be the prior predictive distribution for S^j under the prior $p(\theta|\lambda)$, i.e.

$$p(S^j|\lambda) = \int p(S^j|\theta)p(\theta|\lambda) d\theta.$$

In the definition, the parameter θ in the sampling distribution for S^j given θ is integrated out according to the prior $p(\theta|\lambda)$. $p(S^j|\lambda)$ describes beliefs about S^j before any data are observed for the assumed prior $p(\theta|\lambda)$. Consider the p -values

$$p_{I,b}^j(\lambda) = P(\log p(S^j|\lambda) \leq \log p(h_{I,b}^j|\lambda)), \quad (1)$$

for $S^j \sim p(S^j|\lambda)$ and $j = 1, \dots, k$, $b = 1, \dots, d_I^j$ and

$$p_{P,b}^j(\lambda) = P(\log p(S^j|\lambda) \leq \log p(h_{P,b}^j|\lambda)), \quad (2)$$

where again $S^j \sim p(S^j|\lambda)$ and $j = 1, \dots, k$, $b = 1, \dots, d_P^j$. These p -values give a measure of how far out in the tails of $p(S^j|\lambda)$ the various hypothetical summary values are, and hence how surprising they are. We define a “reasonable” prior $p(\theta|\lambda)$ in the light of the available prior information to be one for which for some appropriate cutoff value α , we have $p_{I,b}^j(\lambda) < \alpha$ for $j = 1, \dots, k$, $b = 1, \dots, d_I^j$ and $p_{P,b}^j(\lambda) \geq \alpha$, $j = 1, \dots, k$, $b = 1, \dots, d_P^j$ (i.e. the values $S^j = h_{I,b}^j$ result in failing a prior predictive check at the cutoff α for $j = 1, \dots, k$, $b = 1, \dots, d_I^j$ and the values $S^j = h_{P,b}^j$, $j = 1, \dots, k$, $b = 1, \dots, d_P^j$ do not fail such a check). Here α is chosen according to the degree of surprise that is considered relevant for

the information we want to put into the prior. It is possible also to use a different cutoff α for different checks. The passing and failing of certain prior predictive checks for hypothetical data summaries represent constraints on what we consider a reasonable prior to be, and we wish to develop methods for searching the hyperparameter space to find corresponding priors satisfying our constraints. The summary statistics can either be univariate or multivariate. However, considering a vector valued S^j is more difficult computationally than considering univariate summaries due to the need to estimate the prior predictive density in (1) and (2) and in our later examples we generally choose univariate S^j . More comments on this, and a cautionary example, are given in Section 5.2. Generally we would want to choose the summary statistics S^j to be reflecting variation related to the parameter θ . This suggests making these summaries sufficient statistics, although non-trivial minimal sufficient statistics do not exist in many problems. Possible choices of the summaries include indicators for the data y belonging to some set (a suggestion made by an anonymous referee), or functions of a point estimator if these are available. Regarding the choice of the hypothetical values, if both plausible and implausible values are specified for a given summary as a pair to convey information about the end point of a plausible range, then making these close together is more constraining. It is important however that the chosen values do not represent information more precise than an expert actually possesses.

As noted in the introduction, while in this work we specify constraints in the form of passing or failing model checks for hypothetical data, the constraints could also be specified in some other way in our procedure, such as through inequalities on quantiles of predictive distributions. The numerical search procedures developed later can also be used with constraints in these other forms. Our method can also apply in situations where prior information is expressed directly on the parameter itself rather than predictively. It is not uncommon for prior distributions to be specified conditionally through a hierarchy, and for marginal prior distributions for functions of the parameter to be unavailable analytically. We

can consider tail probabilities for such marginal priors or inequalities on quantiles for such priors in the same basic framework as our predictive methods. Again, indicator functions for certain sets such as expressing order constraints on certain parameters might be one useful way of adding information.

The p -values (1) and (2) are examples of prior predictive p -values (Box 1980) and such p -values have in particular found use in the checking for prior-data conflicts when the summary statistic is a minimal sufficient statistic (Evans and Moshonov 2006) and for giving a precise formulation of the notion of a weakly informative prior (as in Evans and Jang (2011), inspired by earlier work of Gelman (2006)). When expressing prior information in terms of the results of model checks, the distinction between kinds of checks appropriate for different purposes is related to the choice of summary statistics. This is discussed further in Section 6. In the application here to problems of prior choice it is natural for us to focus on prior predictive checking. However, see also the discussion papers of Gelman et al. (1996) and Bayarri and Berger (2000) or Chapter 5 of Evans (2015) for a variety of perspectives on the broader problem of Bayesian model checking and different types of model checks.

3 History matching and regression ABC methods: An overview

3.1 History matching

History matching is a method used in the literature on assessment of computer models. A computer model or simulator is a complex computer code that takes one or more inputs, which we denote as λ , and produces a set of outputs $\eta(\lambda) = (\eta_1(\lambda), \dots, \eta_k(\lambda))^T$. We are reusing our previous notation for prior hyperparameters deliberately here. In a history match there are some observed data y , intended to correspond to the computer model outputs, and

then a so-called implausibility measure which measures the degree of mismatch between the observations and the computer model output is constructed. The implausibility measure may be based on some implicit or explicit model allowing for measurement error, ensemble variability (the inherent variability of $\eta(\lambda)$ when run multiple times at the same λ when the simulator is stochastic) and model discrepancy (a model term which represents beliefs about lack of fit of the simulator when run at its best input values). In the case of a computationally expensive model, we may also wish to use a flexible interpolator such as a Gaussian process to interpolate or smooth the model outputs $\eta(\lambda)$ based on simulator runs at a limited number of inputs to reduce computational demands. Such a model is called an emulator, and emulation uncertainty at inputs where the computer model has not been run can also be included within the implausibility measure.

History matching proceeds in waves, starting with a space-filling design covering the range of model inputs, and at each wave comes up with a current non-implausible region for the inputs, reducing the size of the non-implausible region at each stage. The phrase non-implausible rather than plausible is used since the non-implausible region consists only of the region of the space not ruled out yet as unsuitable. The iterative aspect of the process allows us to place more points adaptively in “promising” regions of the space, something which is important when λ is high dimensional. If emulation is used for a computationally expensive model, this adaptive aspect where more model evaluations are made in the interesting part of the space allows the quality of emulation to improve as more waves are considered. Thresholds on the implausibility measure determining the current implausible region may become more stringent as the waves proceed and different observations may also be introduced sequentially in this process. The philosophy of history matching is not to find a “best input” for the model, but to explore the space of non-implausible values for the model parameters. The non-implausible region at the end of the process may be empty. A history match can be instructive for guiding model development, and if a model is good

enough to warrant calibration then the history match can be useful for developing efficient computational algorithms. History matching has been successfully used in petroleum reservoir modelling (Craig et al. 1997), galaxy formation models (Vernon et al. 2010; Vernon et al. 2014), rainfall-runoff models (Goldstein et al. 2013), climate models (Williamson et al. 2013) and infectious diseases models (Andrianakis et al. 2015) among other applications. Relationships between history matching and approximate Bayesian computation (ABC) algorithms have been considered recently by Wilkinson (2014) and Holden et al. (2015).

Given an implausibility measure $I(\lambda)$, history matching proceeds in the following way.

1. Initialization. Set $j = 1$ and generate a collection of r points $\lambda_1^{(1)}, \dots, \lambda_r^{(1)}$ for λ according to a space-filling design covering the range of the inputs.
2. Until some stopping rule is satisfied:
 - (a) Calculate $I(\lambda_1^{(j)}), \dots, I(\lambda_r^{(j)})$.
 - (b) Choose some subset of the collection of the current inputs, $\lambda_1^{(j)}, \dots, \lambda_q^{(j)}$, as non-implausible based on thresholding the implausibility measure. This set of points is used to define a current non-implausible region N^j .
 - (c) Generate points $\lambda_1^{(j+1)}, \dots, \lambda_r^{(j+1)}$ according to a new space-filling design covering N^j and set $j = j + 1$.

In Section 4 we describe how we implement the steps in the procedure above for our later applications. There are a variety of approaches in the existing history matching literature for the construction of the implausibility measure, the construction of space filling designs and other choices. In different applications the implausibility measure might change between iterations or only a subset of observations might be considered in the early stages and the implausibility thresholds might change between iterations. In our later applications, at wave j , the wave $j + 1$ samples are generated directly from the current ones without explicitly

defining the set N^j , and so we don't describe how this set is sometimes constructed in the history matching literature. A variety of approaches to this issue may be found in the above references. If an emulator is used in evaluation of the implausibility measure, additional model evaluations could be made at step 2 (b) for the current non-implausible points and the emulator updated appropriately. These additional model evaluations and updating of the emulator may be particularly important in the case of high-dimensional models, and the task of emulation becomes much simpler as the interesting region of the space shrinks over successive waves.

3.2 Regression ABC methods

ABC methods are used in the Bayesian analysis of models where the likelihood is intractable (Tavaré et al. 1997; Pritchard et al. 1999; Beaumont et al. 2002). The basic idea of simple ABC methods is to conduct forward simulations from the model according to parameter values sampled from the prior and to then see whether the simulated data is similar to the observed data. If it is, then the parameter value that generated the simulated data is retained as one that might plausibly have generated the data. A recent review of these methods is given by Marin et al. (2011), but here we confine ourselves to describing only some regression based approaches used in the ABC literature which are relevant to the calculations done in the next section (Beaumont et al. 2002; Blum and François 2010).

Suppose that $p(\theta|\lambda)$ is the prior, $p(y|\theta)$ is the data model and y_{obs} is the observed data. We simulate (θ_i, y_i) , $i = 1, \dots, n$ from the prior and then the simulated data are reduced to a summary statistic $S_i = S(y_i)$ with $S_{obs} = S(y_{obs})$. The role of summary statistics in an ABC analysis is to reduce the dimensionality of the data, and ideally the summary statistics should be nearly sufficient for θ . The idea of regression based ABC methods is to use regression to obtain a conditional density estimate of θ given S_{obs} (i.e. to approximate the posterior distribution $p(\theta|S_{obs})$). We assume that S_{obs} contains most of the relevant information about

θ in y_{obs} . Blum and François (2010), extending methods originally due to Beaumont et al. (2002), consider the regression model

$$\theta_i = \mu(S_i) + \sigma(S_i)\epsilon_i, \quad (3)$$

where $\mu(\cdot)$ and $\sigma(\cdot)$ are flexible mean and standard deviation functions (which they parametrize using neural networks) and the ϵ_i are zero mean variance one residuals. It is assumed above that θ is a scalar parameter, but extensions to the multivariate case are straightforward in which $\mu(S)$ and the ϵ_i are multivariate and $\sigma(S)$ is a matrix square root of the covariance matrix of θ given S . To obtain an approximate sample from $\theta|S_{obs}$, which we write as θ_i^a , $i = 1, \dots, n$ (i.e. an approximate sample from the posterior) we can consider fitting the regression model to obtain estimates $\hat{\mu}(\cdot)$ and $\hat{\sigma}(\cdot)$ of $\mu(\cdot)$ and $\sigma(\cdot)$ respectively, and then use empirical residuals in the fitted regression at $S = S_{obs}$:

$$\theta_i^a = \hat{\mu}(S_{obs}) + \hat{\sigma}(S_{obs})\hat{\epsilon}_i = \hat{\mu}(S_{obs}) + \hat{\sigma}(S_{obs})\hat{\sigma}(S_i)^{-1}(S_i - \hat{\mu}(S_i)),$$

$i = 1, \dots, n$. In the discussion above it is also possible to localize the regression using a kernel function and attach weights to the adjusted sample values θ_i^a .

Nott et al. (2015) consider related methods for repeated conditional density estimation when we want to simulate from a data model for different values of a parameter and where that is expensive. For approximate simulation from the data model the roles of S and θ are reversed in (3). That is, we consider

$$S_i = \mu(\theta_i) + \sigma(\theta_i)\epsilon_i, \quad (4)$$

and then for a given θ an approximate sample from S given θ would be

$$S_i^a = \hat{\mu}(\theta) + \hat{\sigma}(\theta)\hat{\sigma}(\theta_i)^{-1}(\theta_i - \hat{\mu}(\theta_i)),$$

for estimates $\hat{\mu}(\theta)$ and $\hat{\sigma}(\theta)$ of $\mu(\theta)$ and $\sigma(\theta)$. In the next section we use a model similar to (4) to simulate in a computationally thrifty way from a prior predictive distribution $p(S|\lambda)$ for summary statistics S conditional on a prior hyperparameter λ with θ integrated out according to the prior $p(\theta|\lambda)$. Such approximate prior predictive samples are useful for estimating $p(S^j|\lambda)$ (a quantity which appears in our prior predictive p -values (1) and (2)) and hence for choosing an appropriate value of λ .

4 Proposed algorithm for prior choice

Our proposed algorithm applying history matching for prior choice will now be described. If λ denotes prior hyperparameters in a problem of prior choice, given λ we can compute certain features of prior predictive distributions as outputs of the Bayesian model. In the procedure of Section 2 we may consider the outputs to be the p -values in equations (1) and (2). From these an implausibility measure can be constructed based on desired constraints for the outputs. Later we use

$$I(\lambda) = \sum_{j=1}^k \sum_{b=1}^{d_I^j} \max(0, p_{I,b}^j(\lambda) - \alpha) + \sum_{j=1}^k \sum_{b=1}^{d_P^j} \max(0, \alpha - p_{P,b}^j(\lambda)) \quad (5)$$

and we note that $I(\lambda)$ is 0 if the constraints considered in Section 2 are satisfied, i.e. $p_{I,b}^j(\lambda) < \alpha$, $j = 1, \dots, k$, $b = 1, \dots, d_I^j$ and $p_{P,b}^j(\lambda) \geq \alpha$, $j = 1, \dots, k$, $b = 1, \dots, d_P^j$, with $I(\lambda) > 0$ if one or more of these constraints are violated.

Consider once more the expository example of the introduction. There we considered

for the binomial model $\text{Binomial}(n, p)$ parametrized by $\beta = \log p/(1 - p)$ the summary statistic $\hat{p}(1 - \hat{p})/n$ with $\hat{p} = y/n$, and suggested defining some small value of this statistic as implausible as a way of constraining the prior to not place too much mass near values for p of 0 or 1. In this example there is just a single p -value, corresponding to an implausible summary, and the above implausibility measure is given by this p -value minus α if the p -value is bigger than α , and zero otherwise.

The search for prior hyperparameters satisfying the constraints can be performed using the methods of history matching with the implausibility measure (5). One might object that the threshold α used in our implausibility is somewhat artificial. However it should be kept in mind that this threshold is not used in a binary decision making context here, and that the purpose of $I(\lambda)$ is just to guide the search to a fruitful region of the hyperparameter space. Obtaining an exactly 0 value of $I(\lambda)$ may not be so important. The use of p -values in $I(\lambda)$ is convenient for the way that it puts information from the different summary statistics on the same scale, and we have found the choice (5) for the implausibility measure to be useful although there are certainly other ways that the implausibility could be defined.

Steps 2 b) and c) of the history matching algorithm given in Section 3.1 for wave j are implemented in our later examples in the following way. First, choose some fraction γ of r in such a way that both $1/\gamma$ and $q = \gamma r$ are integers. For instance, in the first example of Section 5 we use $\gamma = 0.1$ and $r = 100$. Next, choose the q values of λ in the current wave for which $I(\lambda)$ is smallest. Write these values as $\lambda_1^{*(j)}, \dots, \lambda_q^{*(j)}$. Then for each of $k = 1, \dots, q$, generate $1/\gamma$ values from a normal distribution $N(\lambda_k^{*(j)}, \Sigma^{(j)})$ where $\Sigma^{(j)} = h^2 V_j$, V_j is the sample covariance matrix of all the wave j samples, and $h = \left(\frac{4}{(2d+1)q} \right)^{1/(d+4)}$ where d is the dimension of λ . Note that this results in $q/\gamma = r$ samples that we take as the wave $j + 1$ samples. In our later examples we use the modified sampling approach in the `mvrnorm` function in the R package `MASS` (Venables and Ripley 2002) with the option `empirical=TRUE` to obtain generated samples that have exactly the sample covariance matrix $\Sigma^{(j)}$. The definition of $\Sigma^{(j)}$

in the sample generation step is obtained by inflating a conventional choice of kernel used in multivariate kernel density estimation by a factor of 4 (Silverman 1986). There are other ways to generate a space-filling design for each wave; the idea above and that we implement later in examples is a simple one based on a similar suggestion in Andrianakis et al. (2015) based on perturbing values according to a normal kernel with enough variability to ensure that the new points are sufficiently different to the current one. The intuition behind our choice for h is that after pruning away the implausible samples in the current wave, we want to generate a set of points for the next wave that covers the distribution for the current set of non-implausible samples. The kernel estimate with inflation of the conventional bandwidth is just to make the next wave samples somewhat overdispersed compared to the distribution of current non-implausible samples. It is difficult to say anything about optimality of our suggested choice of h . A larger value of h will ensure that the non-implausible region is not collapsed down too quickly, at the expense of additional computations. How quickly we should narrow down the non-implausible region also interacts with how many samples are used in the initial space-filling design, and how smooth the implausibility measure is. The only remaining detail to specify in the algorithm is the stopping rule. We don't specify a formal rule for this but instead examine the results of the latest wave graphically every few iterations to see whether the implausible region continues to be reduced or an acceptable match to the constraints has been found.

Computing the implausibility measures in the application of history matching to prior choice as discussed in Section 3 involves computation of the p -values in equations (1) and (2) for a large number of different values of λ and this can be computationally burdensome. Our solution is to use the regression approximate Bayesian computation (ABC) methods introduced in Section 3.2 to approximate these p -values in a computationally thrifty way. The methods considered are based on those developed in Nott et al. (2015), and play a similar role in our later examples to the use of emulators in history matching for computationally

expensive computer models.

Suppose we wish to approximate $I(\lambda)$ for a possibly large set of different λ values, λ^i , $i = 1, \dots, m$. These values might be a grid over the region of interest for λ if λ is low-dimensional, or in the history matching procedure they might be the hyperparameter values generated in the current wave. Let $p(\lambda)$ be a pseudo-prior for λ which covers the range of the values of λ of interest. This pseudo-prior is not to be used for inference but is used in generation of samples of the summaries S^j . We simulate values $(\lambda_i, \theta_i, y_i)$ from $p(\lambda)p(\theta|\lambda)p(y|\theta)$, $i = 1, \dots, n$ independently. From the y_i we obtain simulated summaries $S_i^j = S^j(y_i)$, $i = 1, \dots, n$, $j = 1, \dots, k$. We can obtain an approximate sample from $p(S^j|\lambda)$ for any given value of λ by considering the regression adjustment methods of Section 3 applied to the regression model

$$S_i^j = \mu^j(\lambda_i) + \sigma^j(\lambda_i)\epsilon_i,$$

where the ϵ_i are independent and identically distributed errors with variance one and $\mu^j(\lambda)$ and $\sigma^j(\lambda)$ are flexible mean and standard deviation functions. This is similar to the regression adjustment approach considered for equation (4) in Section 3 applied to the marginalized model for the summaries where θ has been integrated out according to $p(\theta|\lambda)$. Extension to the case where S_i^j is multivariate can also be considered but in our later examples the S^j are each univariate summaries. Fitting the regression model locally, based on a certain number of nearest neighbours of λ^l , is often useful. This is something we consider later in the examples with a nearest neighbour distance following the default choice in the **R** package **abc** (Csilléry et al. 2012).

An approximate sample from $p(S^j|\lambda^l)$ is

$$S_i^{j,l} = \hat{\mu}^j(\lambda^l) + \hat{\sigma}^j(\lambda^l)\hat{\sigma}^j(\lambda_i)^{-1}(S_i^j - \hat{\mu}^j(\lambda_i)), \quad i = 1, \dots, n \quad (6)$$

and then we can construct a kernel density estimate of $p(S^j|\lambda^l)$, written $\hat{p}(S^j|\lambda^l)$, from these approximate samples. The computation of the estimated p -values $\hat{p}_{I,b}^j(\lambda^l)$, $j = 1, \dots, k$, $b = 1, \dots, d_I^j$ and $\hat{p}_{P,b}^j$, $j = 1, \dots, k$, $b = 1, \dots, d_P^j$, can be performed using the following algorithm.

1. Given the input hyperparameter λ^l , obtain approximate samples $S_i^{j,l}$, $i = 1, \dots, n$ from $p(S^j|\lambda^l)$, $j = 1, \dots, k$, according to (6).
2. For each statistic S^j , $j = 1, \dots, k$, calculate a kernel estimate of $p(S^j|\lambda^l)$ at $S_i^{j,l}$, $\hat{p}(S_i^{j,l}|\lambda^l)$, $i = 1, \dots, n$, $h_{I,b}^j$, $b = 1, \dots, d_I^j$ and $h_{P,b}^j$, $b = 1, \dots, d_P^j$.
3. Calculate

$$\hat{p}_{I,b}^j(\lambda^l) = \frac{1}{n} \sum_{i=1}^n I(\log \hat{p}(S_i^{j,l}|\lambda^l) \leq \log \hat{p}(H_{I,b}^j|\lambda^l)), \quad j = 1, \dots, k, b = 1, \dots, d_I^j$$

and

$$\hat{p}_{P,b}^j(\lambda^l) = \frac{1}{n} \sum_{i=1}^n I(\log \hat{p}(S_i^{j,l}|\lambda^l) \leq \log \hat{p}(H_{P,b}^j|\lambda^l)), \quad j = 1, \dots, k, b = 1, \dots, d_P^j$$

Given the estimated p -values for a certain λ^l we can check whether λ^l is acceptable according to our criteria by checking if $\hat{p}_{I,b}^j(\lambda^l) < \alpha$, $j = 1, \dots, k$, $b = 1, \dots, d_I^j$ and $\hat{p}_j(\lambda^l) \geq \alpha$, $j = 1, \dots, k$, $b = 1, \dots, d_P^j$. An approximate implausibility value $I(\lambda^l)$ can also be computed from the p -values. Note that the regression ABC computations are screening computations, and high accuracy is not needed. Once a hyperparameter value is chosen based on the regression calculations as giving a prior satisfying the desired constraints we can check its suitability. We can do this by generating a large number of values of S^j , $j = 1, \dots, k$ from the prior predictive distribution for the chosen λ , and from these approximate the p -values accurately, to check that the regression approximations were good enough. Alternatively,

we can dispense with p -values and just look at plots of prior predictive densities for the summary statistics directly.

The approach we have described of approximating prior predictive samples based on local regression adjustments can fail when the prior predictive density changes rapidly as a function of λ , and it may also be difficult to apply in high dimensions. It is also assumed above that summary statistics are generated once at the beginning of the history match according to values for λ simulated under the pseudo-prior $p(\lambda)$. It was mentioned in Section 3 that a powerful aspect of history matching is the way that additional model evaluations (or summary statistic simulations in the present case) can be made as the waves of the history matching proceed. That is, we can generate additional summary statistic simulations at each of the current non-implausible λ values in the history matching waves to improve the quality of the regression adjustment approach for approximating the prior predictive distribution in the interesting parts of the hyperparameter space. This is most interesting when the number of hyperparameters is high-dimensional, and for our highest dimensional example later (with four hyperparameters) we consider such an approach. Emulation methods are thoroughly developed in the existing literature for deterministic computer models. But, where stochastic models are considered and the task is to emulate the distribution of an output as a function of inputs, simple methods such as just emulating means and variances are often considered, depending on what is required for the chosen implausibility measure. In our application, capturing more complex features of the prior predictive density becomes important. The regression ABC approach outlined here is not the only one that could be considered. However, a comparison of different conditional density estimation methods in this application is beyond the scope of the present work.

5 Examples

We illustrate our methodology in three examples. In the first two examples there are just two hyperparameters to be chosen and we can plot the way that the predictive p -values in our checks vary with the hyperparameters over a grid; such plots are useful for checking the results of the history match. Both the p -values at the grid points in these plots, as well as the p -values used to approximate the implausibility measure for the history matching samples, are obtained using regression ABC approximations to the prior predictive densities of the summaries. In the third example there are four hyperparameters to be chosen, and consideration of a grid of hyperparameter values is no longer feasible.

5.1 Logistic regression example

We consider a logistic regression for an experiment described in Racine et al. (1986) where 5 animals at each of 4 dose levels were exposed to a toxin. We write the dose levels as $x_1 < x_2 < x_3 < x_4$ and assume that these values have been transformed to a log scale, centered and scaled as in Gelman et al. (2008). If y_i is the number of animals killed at dose level x_i , the data model is $y_i \sim \text{Binomial}(5, p_i)$ with $\log p_i / (1 - p_i) = \beta_0 + \beta_1 x_i$. Gelman et al. (2008) consider a prior on β where β_0 and β_1 follow independent Cauchy distributions centered on zero with scale $\lambda_1 = 10$ and $\lambda_2 = 2.5$ respectively. Here we consider $\lambda = (\lambda_1, \lambda_2)$ as hyperparameters to be chosen, with $\lambda \in [0.5, 10] \times [0.5, 10]$.

Our elicitation method requires us to specify some hypothetical data to be plausible or implausible under the prior. Write $\hat{\beta} = (\hat{\beta}_0, \hat{\beta}_1)$ for the posterior mode of β based on independent normal $N(0, 100)$ priors on β_0, β_1 . Note that $\hat{\beta}$ is similar to the MLE in non-degenerate settings but will exist even when the MLE does not. For each dose x_i , let $\hat{p}_i = 1 / (1 + \exp(-\hat{\beta}_0 - \hat{\beta}_1 x_i))$ be the corresponding fitted probability of death at dose x_i under the fitted model. Let us consider the summary statistic $S^1 = \sum_{i=1}^4 5\hat{p}_i(1 - \hat{p}_i)$ which

is the sum of the variances of the responses when $\beta = \hat{\beta}$. The statistic S^1 will tend to be small if all the responses are close to either zero or the maximum value of 5 resulting in fitted probabilities at the different dose levels all close to zero or one. If all \hat{p}_i are equal to either 0.01 or 0.99, then the value of S^1 would be 0.198 and we might wish the prior to express the information that this is an implausible value for S^1 . S^1 defined here is the natural extension to the logistic regression case of the summary statistic used in the expository example of the introduction.

In this example we might also expect that it would not be surprising if the fitted probability of death goes from a value near zero at the lowest dose to a value near 1 at the highest dose, in a fairly smooth way. If $\hat{p}_1 = 0.01$, $\hat{p}_2 = 0.25$, $\hat{p}_3 = 0.75$ and $\hat{p}_4 = 0.99$, then the corresponding value of S^1 would be 1.974. We consider a prior within our framework in which $S^1 = 0.198$ is considered to be implausible, and $S^1 = 1.974$ is considered to be plausible. This is weak prior information, but enough to constrain hyperparameter choice in a useful way. S^1 is treated as a continuous quantity in our calculations, even though it is discrete. This is a reasonable approximation when the number of different possible values is large, as here.

For the hypothetical data summary $S^1 = 0.198$, we compute the predictive p -value for the summary statistics chosen using the method of Section 4 and using a grid of 10,000 λ values in our target range $\lambda \in [0.5, 10] \times [0.5, 10]$ with the grid formed from 100 equally spaced values in each dimension. The regression adjustment calculations for computation of the p -values are done using the default implementation of the `abc` function in the `abc` R package Csilléry *et al.* (2012). We used 400,000 simulated values of the summary statistics (S^1, S^2) , local linear regression adjustment and 1,000 nearest neighbours in the localized regression ABC procedure. A plot of how the p -value changes as a function of λ is shown in the left panel of Figure 1. Note the two blue regions in the graph where the p -value is small; the region on the left occurs for hyperparameter values where 0.198 is an implausibly small value, whereas




Figure 1: Conflict p -value as a function of λ for logistic regression example. p -value for check for $S^1 = 0.198$ (left) and for $S^1 = 1.974$ (right). In both graphs the overlaid points are from the third wave of the history match and the minimum implausibility obtained is zero.

the region on the right occurs for hyperparameter values for which 0.198 is implausibly large. A similar plot of the p -value as a function of λ for the check with $S^1 = 1.974$ is shown in the right panel. An acceptable value for λ is a value in the light blue region in the left panel (small p -value indicating a prior-data conflict) and avoiding the light blue region in the right panel (a p -value which is not small indicating the absence of a conflict). The points overlaid on the graphs are obtained from using the history matching method of Section 3. In the history match the algorithm is initialized with a maximin latin hypercube design of $r = 100$ points, $\gamma = 0.1$ and the points shown in the graph are the retained values after 4 waves. The p -values in the implausibility measure are again computed using the method of Section 4. The minimum implausibility obtained is 0, i.e. we are successful at finding hyperparameter values satisfying the constraints. As mentioned above, in considering this example Gelman et al. (2008) considered a default prior with $\lambda_1 = 10$ and $\lambda_2 = 2.5$. This is a weakly informative choice, and it can be seen from Figure 1 that to match the information we have suggested putting into our analysis a smaller value of λ_1 is needed. Also shown in Figure 2 are the marginal posterior distributions of β_0 and β_1 for the default prior with $\lambda_1 = 10$ and




Figure 2: Marginal posterior distributions for β_0 (left) and β_1 (right) for default prior with $(\lambda_1, \lambda_2) = (10, 2.5)$ as well as history matching hyperparameter values of $(\lambda_1, \lambda_2) = (0.33, 2.08)$ and $(\lambda_1, \lambda_2) = (0.23, 0.73)$ (labelled “HistMatchABC1” and “HistMatchABC2” respectively).

$\lambda_2 = 2.5$, as well as for two hyperparameter values obtained from the history match. The posterior distributions are computed for the observed data of $(y_1, y_2, y_3, y_4) = (0, 1, 3, 5)$. In this example it is seen that the prior information we have put in makes some difference to the resulting inference, particularly for the intercept.

5.2 Sparse signal shrinkage prior

Next we consider prior choice for a linear model with a sparse signal shrinkage prior on the coefficients. The shrinkage prior we consider is the horseshoe+ prior of Bhadra et al. (2015). The need in modern data analysis to consider increasingly complex models with respect to both the number of parameters and hierarchical structure has resulted in a very large literature on sophisticated shrinkage priors in a range of applications. We consider only the horseshoe+ prior for a high-dimensional linear model in this example, but the kind of analysis we do here could be done for other shrinkage priors, of which there are many. Bhadra et al.

(2015) give a survey of the current state of the art in the area. We describe a general version of our model first which also incorporates observation specific mean shift terms that can account for outliers in the model, using similar ideas to those considered in She and Owen (2011). A simplified version of the model with two hyperparameters will be considered in this subsection, and the more general form of the model with four hyperparameters will be considered in the next subsection.

For some design matrix X ($n \times p$) consider the model

$$y = \beta_0 1_n + X\beta + \delta + \epsilon, \quad (7)$$

where $y = (y_1, \dots, y_n)^T$ is an n -vector of responses, β_0 is an intercept term, 1_n denotes an n -vector of ones, β is a $p \times 1$ vector of regression coefficients, $\delta = (\delta_1, \dots, \delta_n)^T$ is an n -vector of mean shift parameters intended to be sparse and which allows for outliers in a small number of observations, and $\epsilon \sim N(0, \sigma^2 I)$ are zero mean normal independent and identically distributed residuals. The model is not identifiable unless sparsity assumptions are made for δ , and in the case where $p > n$, which is the case we consider here, we also need to make some assumptions of sparsity for β .

We consider a Bayesian analysis with priors $\beta_0 \sim N(0, \sigma_0^2)$ and $\sigma \sim \text{HC}(0, A_\sigma)$ (where $\text{HC}(0, A_\sigma)$ denotes the half Cauchy distribution with scale parameter A_σ). The elements of β are independent in their prior, $\beta_j \sim N(0, \sigma_j^2)$, with $\sigma_j \sim \text{HC}(0, A_\beta \gamma_j)$, $\gamma_j \sim \text{HC}(0, 1)$ and A_β is a scale parameter to be chosen. Similarly in the prior for δ the elements of δ are independent in the prior with $\delta_j \sim N(0, \tau_j^2)$, $\tau_j^2 \sim \text{HC}(0, A_\delta \zeta_j)$, $\zeta_j \sim \text{HC}(0, 1)$ where A_δ is a hyperparameter to be chosen. The prior specification is complete once the hyperparameters σ_0^2 , A_σ , A_β and A_δ are fixed. In the current section we consider the model where $\delta = 0$ and hence there is no need to set A_δ and where σ_0^2 is fixed at 100. The full model is considered further in the next subsection.

We consider choice of (A_σ, A_β) in the context of the sugar data set considered in Brown et al. (1998). In this dataset there are $p = 700$ predictors in the training sample, 3 response variables and 125 observations in the training set, so that we are considering a case where $p > n$. We consider the response variable glucose and center and scale all columns of the design matrix. Now consider applying our method. For summary statistics, we define S^1 to be the log of the marginal variance of y averaging over the predictors, i.e. $S^1 = \log s^2$ where

$$s^2 = \frac{1}{n-1} \sum_{j=1}^n (y_j - \bar{y})^2,$$

where \bar{y} is the sample mean of y . Some idea of the range of the responses marginally is very likely to be available in applications and so it may be easy to specify what would be plausible or implausible values for S^1 . We consider $S^1 = \log 16$ to be plausible and $S^1 = \log 50$ to be implausible (the marginal variance for the observed data is about 16 here).

We also consider another summary statistic $S^2 = S^2(y)$ defined as follows. The summary statistic is an adjusted R^2 type measure of how much variation is explained by the predictors, but one that is appropriate to the situation of more covariates than observations and which is based on a simple version of the refitted cross-validation method of Fan et al. (2012). Specifically, for any y , we split the data into 2 halves randomly, rank the predictors by absolute correlation with the response on one half of the data, fit to the other half of the data using only the $n/4$ highest ranked variables on the first half, obtain the adjusted R^2 value, and then swap the roles of the two halves and average the two adjusted R^2 estimates obtained. This gives an adjusted R^2 estimate for a certain random split. We split randomly 10 times and then our summary statistics $S^2(y)$ is the average of the adjusted R^2 estimates over the ten splits. We want to require that both $S^2 = 0.05$ as well as $S^2 = 0.95$ are plausible, so that the model allows both a small or large amount of variation in the response variable to be explainable through the regression *a priori*.

Figure 3 shows plots of the p -values for the tests based on the four summary statistics as (A_σ, A_β) vary. The plots are for 100×100 grids equally spaced in each dimension for $(A_\sigma, \log A_\beta)$ covering the range $[0, 2] \times [-\log 100p, -\log p]$. The regression adjustment calculations for computation of the p -values are done using 100,000 summary statistics with local linear regression adjustments and 1,000 nearest neighbours. Similar to the last example overlaid on the graphs are the retained points from the third wave of a history match implemented in the same way as the previous example with $r = 100$ and $\gamma = 0.1$. The history match succeeds in finding prior hyperparameter values corresponding to priors which satisfy the constraints. In the top right plot we want to be in the darkest blue region (i.e. the corresponding summary is implausible), and in the other plots we want to avoid the darkest blue region (i.e. the corresponding summaries are plausible). Note that avoiding the darkest blue region here means the p -values should be larger than a 0.05 cutoff, which should be interpreted according to the colour scale shown at the right of the plots and the points can still be in the blue (but not the darkest blue) region.

It is interesting to see what happens in this example when we change the prior on β to $\beta_j \sim N(0, A_\beta)$, so that now A_β is a scale parameter to be chosen in a normal prior, but where our predictive constraints remain the same. We continue to use the notation A_β for the scale parameter in the prior on β even though this is of course a different parameter in the two priors. State of the art sparsity inducing priors like the horseshoe+ have good frequentist performance in a number of senses as described in Bhadra et al. (2015). Here we illustrate a more Bayesian way in which this prior is good in this example. Before we did a history match in this example we expected that the normal prior would work poorly in the sense of not being able to capture the information that either a large or small amount of the variation in the response should be explainable through the covariates *a priori*. Our intuition was incorrect, and it was in fact possible to satisfy our constraints. The results of wave 5 of our history match for the normal prior are shown in Figure 4.




Figure 3: Conflict p -value as a function of (A_σ, A_β) for sparse signal shrinkage example. p -value for check for $S^1 = \log 16$ (top left), $S^1 = \log 50$ (top right), $S^2 = 0.05$ (bottom left) and $S^2 = 0.95$ (bottom right). In both graphs the overlaid points are from the third wave of the history match and the minimum implausibility obtained is zero.




Figure 4: Conflict p -value as a function of (A_σ, A_β) for normal prior example. p -value for check for $S^1 = \log 16$ (top left), $S^1 = \log 50$ (top right), $S^2 = 0.05$ (bottom left) and $S^2 = 0.95$ (bottom right). In both graphs the overlaid points are from the third wave of the history match and the minimum implausibility obtained is 0.

However, now consider the following. If $S^1 = \log 16$ and $S^2 = 0.95$ should both be plausible, perhaps we should also require that $(S^1, S^2) = (\log 16, 0.95)$ should be plausible in the joint prior predictive for (S^1, S^2) . Figure 5 shows kernel estimates of the joint prior predictive density for (S^1, S^2) for the horseshoe+ and normal priors for two particular hyperparameter values achieving zero implausibility, based on 1000 prior predictive samples. We can see that $(S^1, S^2) = (\log 16, 0.95)$ is plausible for the horseshoe+ prior, but not for the normal prior. The explanation for this is that it is only when the noise variance is small that the regression can explain a lot of the variation in the case of the normal prior. The behaviour of the horseshoe+ prior, however, is more acceptable. This example illustrates perhaps some of the pitfalls of considering plausible and implausible values for one-dimensional summary statistics separately. While this is a useful strategy for defining constraints, and it makes computations more convenient, once a reasonable candidate hyperparameter value is found it may be useful to consider the behaviour of the joint prior predictive for several summaries simultaneously.

5.3 An example with higher-dimensional hyperparameter

Continuing the last example, consider the full model (7) described in Section 5.2 where now we allow δ to be nonzero. We also consider the situation where σ_0^2 is not fixed in the prior for β_0 . Now we have four hyperparameters to be chosen, $(\sigma_0, A_\sigma, A_\beta, A_\delta)$. Unlike the previous two examples with only two hyperparameters, it is not feasible to use a grid-based approach to produce plots of how the conflict p -values vary over the hyperparameters for comparison with the results of the history match. We retain the summary statistics and constraints of Section 5.2, with the difference that s^2 is replaced by a robust measure of scale (the median absolute deviation estimator), and in the linear regression fits for the refitted cross-validation procedure we use the robust `lmrob` function in R (Rousseeuw et al. 2015) to obtain the adjusted R^2 estimate. We also add to the constraints of Section 5.2 three




Figure 5: Prior predictive densities for (S^1, S^2) for two zero implausibility hyperparameter values for horseshoe+ prior (left) and normal prior (right). The point $(S^1, S^2) = (\log 16, 0.95)$ is marked. The hyperparameters are $(A_\sigma, A_\beta) = (0.36, 0.014)$ for the normal prior, and $(A_\sigma, A_\beta) = (0.033, 0.00004)$ for the horseshoe+ prior.

additional constraints. We choose a summary statistics S^3 to be the log of the absolute value of the median of the responses, and specify $S^3 = \log 15$ to be plausible, and $S^3 = \log 20$ to be implausible. As an additional summary statistic we use the following procedure. We consider the log sample kurtosis of the residuals obtained from the `lmrob` function averaged over 10 split samples using the same refitted cross-validation procedure as for the adjusted R^2 measure. This is intended to be some sample measure of the “tailedness” of the distribution. Writing S^4 for this statistic, we consider $S^4 = \log 50$ to be implausible. The value of $\log 50$ was obtained as the log of the approximate median of sample kurtosis values from a Cauchy distribution sample of size 125. Note that we use sample kurtosis here as a summary of the data without worrying about whether any corresponding population quantity exists. The information in this last summary statistic is intended to state the requirement that we should not have a very large proportion of very extreme outliers. Figure 6 shows pairwise scatter plots of the hyperparameter values on a log scale in wave 1 through wave 5 of a history match

with $r = 1000$ and $\gamma = 0.1$ and the first wave initialized with a maximin latin hypercube design covering the range $[e^{-3}, e^2] \times [e^{-5}, e] \times [10^{-6}, 0.5] \times [10^{-6}, 0.5]$ for the hyperparameters. The history match succeeds in finding prior hyperparameter values corresponding to priors which satisfy the constraints.

In Section 4 it was mentioned that it may be helpful to adaptively generate new summary statistic simulations as the waves of the history match proceed. The results of Figure 6 were obtained without doing this, using 100,000 simulations at the beginning of the procedure. Figure 7 shows 8 waves of a history match where the initial number of summary statistic simulations was reduced to 10,000, with 1,000 additional simulations added at each wave (100 further simulations at each of the 10 non-implausible values retained at each wave). The results are similar to before, but now the total number of model simulations has been reduced to 18,000 rather than 100,000. Although this is not a very high-dimensional example, this illustrates the point that this adaptive approach to the model simulations to improve the quality of the regression ABC adjustment can be very important as the number of hyperparameters increases. Effectively the additional model simulations allow us to use smaller neighbourhoods in this local nonparametric procedure. Any approach to flexible conditional density estimation could be used instead of the regression ABC approach for approximating the prior predictive densities as a function of the hyperparameters, but any such alternative method will also benefit from additional model simulations in the important region of the space. Figure 8 shows estimated prior predictive densities of the summary statistics used in the history match obtained from one of the hyperparameter values with implausibility zero in Figure 6, $(\sigma_0, A_\sigma, A_\beta, A_\gamma) = (3.91, 0.016, 0.000013, 0.000045)$, with implausibility measure 0. The graphs presented are histograms and kernel estimates based on 1000 prior predictive samples.

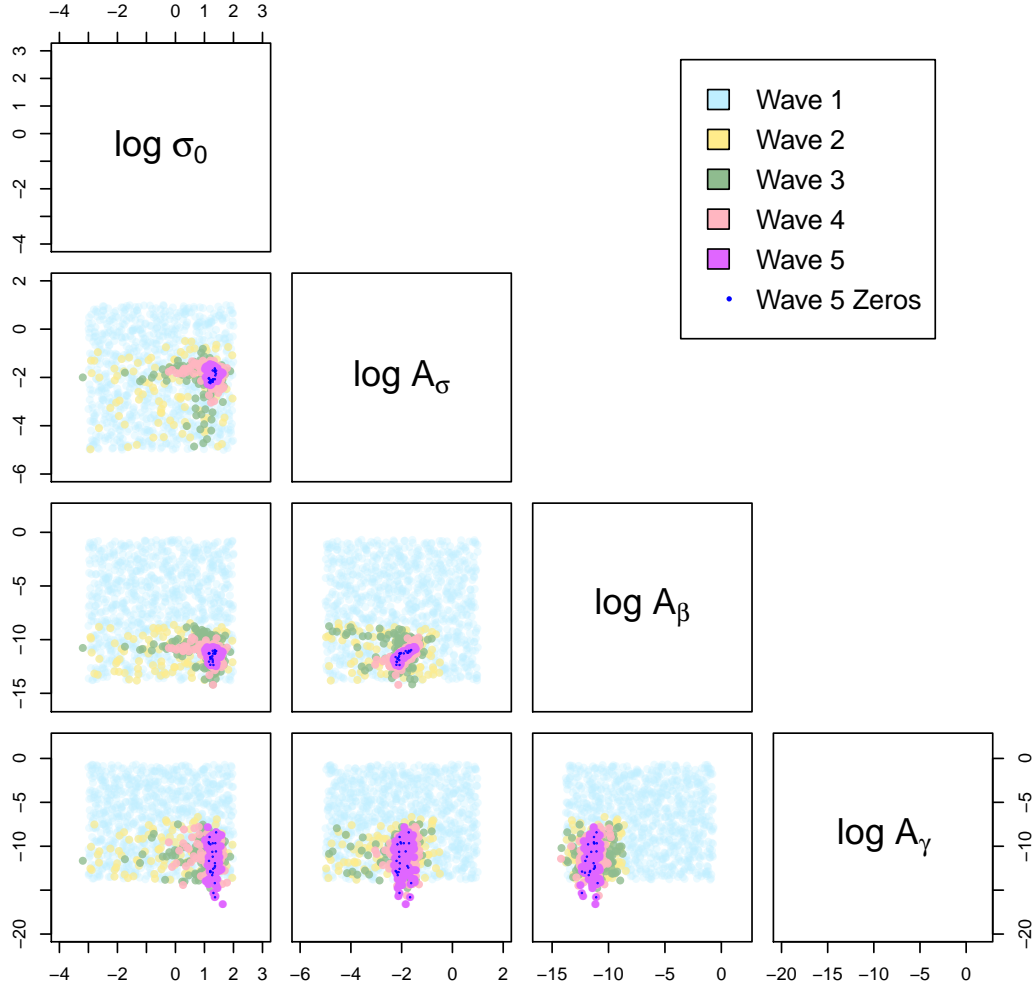


Figure 6: Pairwise scatterplots of hyperparameters on log scale of wave 1 to wave 5 of the history match. The minimum implausibility value obtained in wave 5 is 0.

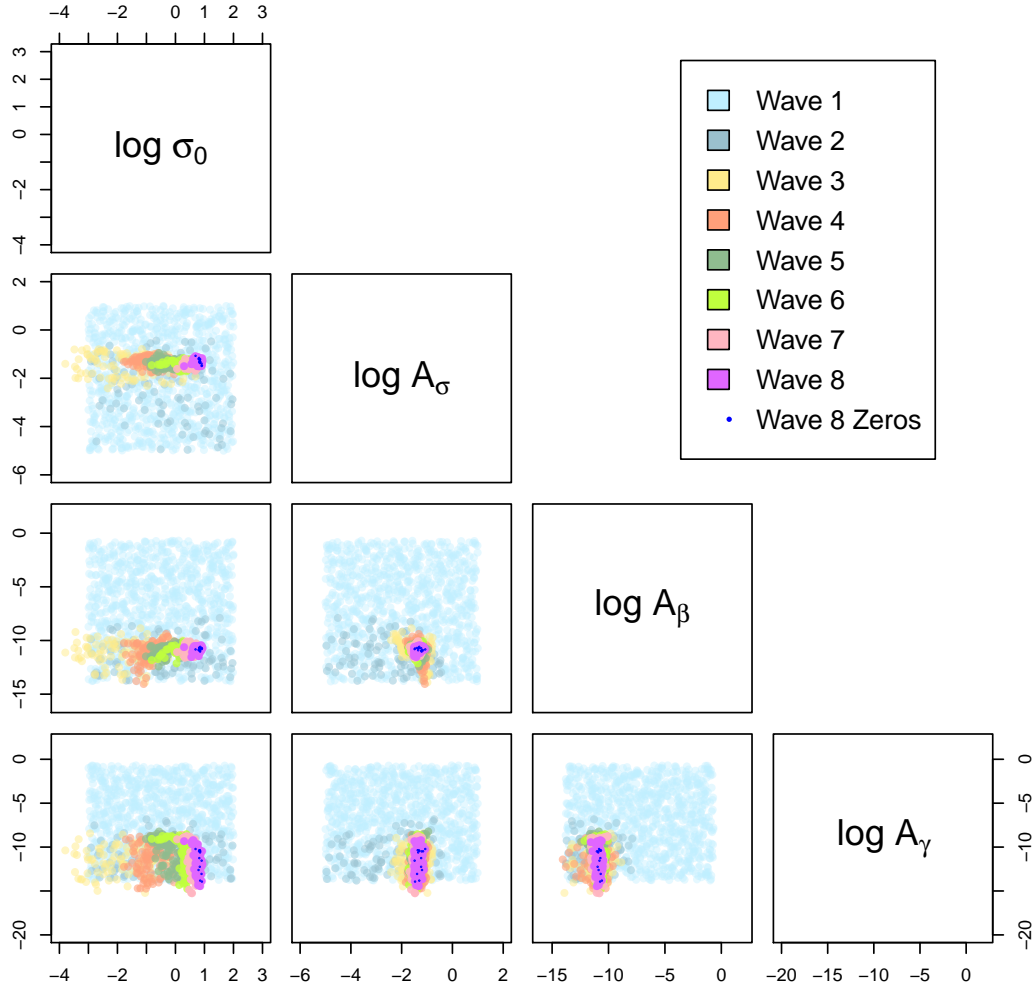


Figure 7: Pairwise scatterplots of hyperparameters on log scale of wave 1 to wave 8 of the history match with additional model simulations at each wave. The minimum implausibility value obtained in wave 8 is 0.

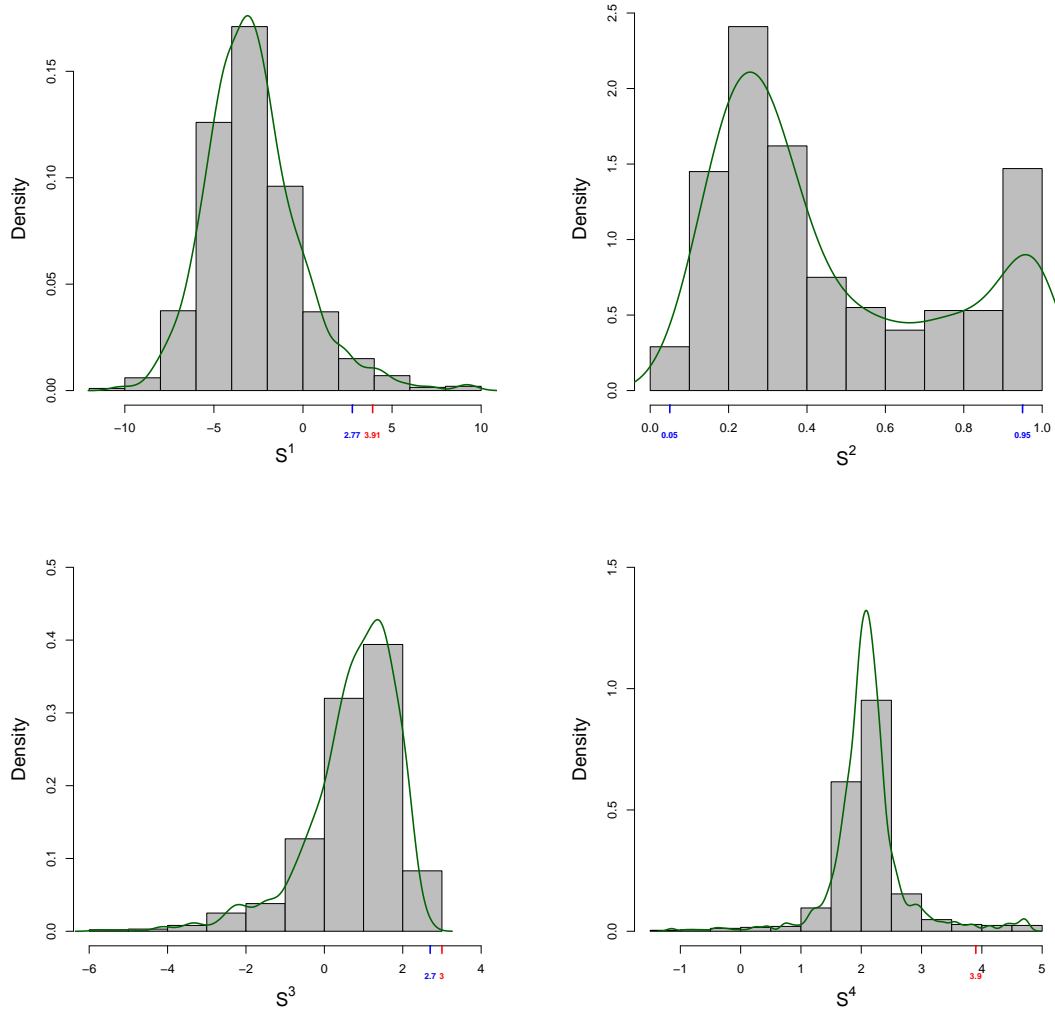


Figure 8: Prior predictive densities of S^1 , S^2 , S^3 , S^4 for hyperparameter value achieving zero implausibility.

6 Discussion

We have considered a novel application of the ideas of history matching used in the assessment of computer models to the problem of prior choice. By defining the implausibility measure in the history match through some prior predictive constraints, we are able to implement predictive elicitation even for complex models. Regression adjustment ABC methods are also used to ease the computational burden in application of the method. We believe the analyses presented in some of the examples are insightful, and in some cases led to some new understanding of the effects of the parameter prior on the prior predictive densities.

Further investigation is needed to see how well the methods we have developed scale to problems where the number of hyperparameters is much larger. Also, it is not clear whether the specific form for the implausibility measure that was chosen was the best one, or whether the various tuning parameter choices in the history matching procedure are best or capable of a more automated implementation. Although, as we have stressed throughout the manuscript, we are focusing mostly on computational questions in this paper it is also worth considering how the methods and algorithms developed are best integrated within an elicitation procedure in complex applied problems.

In the literature on Bayesian model checking a distinction is sometimes made between checking the data model and checking for prior-data conflict, and in hierarchical models one can also try to check the different levels of the prior. It would be interesting to consider more carefully the roles of different kinds of checks in specifying prior predictive constraints and whether that might be informative *a priori* about the possible need to change the data model rather than the parameter prior. Also in the hierarchical setting if checks are considered for different levels of the hierarchical prior perhaps that can be informative of how a hierarchical prior might need to be changed to satisfy certain predictive constraints in the case where a suitable prior is not found in a preliminary analysis. The ABC computations in our method

are similar to those used in Nott et al. (2015) for finding weakly informative priors and many of the elicitation calculations can be reused for finding such a weakly informative prior in the event that there is a prior-data conflict. Also worthy of further investigation is whether greater use can be made of the full set of prior distributions returned by the history match. Here we have simply focused on choice of a single “adequate” prior but there is a richer source of information that can be used in the results of the history matching procedure.

Acknowledgements

David Nott was supported by a Singapore Ministry of Education Academic Research Fund Tier 2 grant (R-155-000-143-112). Christopher C Drovandi was supported by an Australian Research Councils Discovery Early Career Researcher Award funding scheme (DE160100741). Kerrie Mengersen was supported by an Australian Research Council Laureate Fellowship.

References

- Andrianakis, I., I. R. Vernon, N. McCreesh, T. J. McKinley, J. E. Oakley, R. N. Nsubuga, M. Goldstein, and R. G. White (2015). Bayesian history matching of complex infectious disease models using emulation: A tutorial and a case study on HIV in Uganda. *PLOS Computational Biology* 11(1), e1003968+.
- Bayarri, M. J. and J. O. Berger (2000). P values for composite null models (with discussion). *Journal of the American Statistical Association* 95, 1127–1142.
- Beaumont, M. A., W. Zhang, and D. J. Balding (2002). Approximate Bayesian computation in population genetics. *Genetics* 162, 2025–2035.
- Bedrick, E. J., R. Christensen, and W. Johnson (1996). A new perspective on priors for generalized linear models. *Journal of the American Statistical Association* 91, 1450–

1460.

- Bhadra, A., J. Datta, N. G. Polson, and B. Willard (2015). The Horseshoe+ estimator of ultra-sparse signals. arXiv1502.00560.
- Blum, M. G. B. and O. François (2010). Non-linear regression models for approximate Bayesian computation. *Statistics and Computing* 20, 63–75.
- Box, G. E. P. (1980). Sampling and Bayes’ inference in scientific modelling and robustness (with discussion). *Journal of the Royal Statistical Society, Series A* 143, 383–430.
- Brown, P. J., M. Vannucci, and T. Fearn (1998). Multivariate Bayesian variable selection and prediction. *Journal of the Royal Statistical Society, Series B* 60, 627–641.
- Craig, P. S., M. Goldstein, A. H. Seheult, and J. A. Smith (1997). Pressure matching for hydrocarbon reservoirs: A case study in the use of Bayes linear strategies for large computer experiments. In C. Gatsonis, J. S. Hodges, R. E. Kass, R. McCulloch, P. Rossi, and N. D. Singpurwalla (Eds.), *Case Studies in Bayesian Statistics*, Volume 121 of *Lecture Notes in Statistics*, pp. 37–93. Springer New York.
- Csilléry, K., O. François, and M. G. B. Blum (2012). ABC: an R package for approximate Bayesian computation (ABC). *Methods in Ecology and Evolution* 3, 475–479.
- Daneshkhah, A. and J. Oakley (2010). Eliciting multivariate probability distributions. In *Rethinking risk measurement and reporting, Volume 1* (Ed. K. Böcker). Böcker, K., Risk Books, London.
- Evans, M. (2015). *Measuring Statistical Evidence Using Relative Belief*. Taylor & Francis.
- Evans, M. and G. H. Jang (2011). Weak informativity and the information in one prior relative to another. *Statistical Science* 26, 423–439.
- Evans, M. and H. Moshonov (2006). Checking for prior-data conflict. *Bayesian Analysis* 1, 893–914.

- Fan, J., S. Guo, and N. Hao (2012). Variance estimation using refitted cross-validation in ultrahigh dimensional regression. *Journal of the Royal Statistical Society, Series B* 74, 37–65.
- Garthwaite, P. H. and J. M. Dickey (1988). Quantifying expert opinion in linear regression problems. *Journal of the Royal Statistical Society, Series B* 50, pp. 462–474.
- Garthwaite, P. H., J. B. Kadane, and O. Anthony (2005). Statistical methods for eliciting probability distributions. *Journal of the American Statistical Association* 100, 680–701.
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper). *Bayesian Analysis* 1, 515–534.
- Gelman, A., A. Jakulin, M. G. Pittau, and Y.-S. Su (2008). A weakly informative default prior distribution for logistic and other regression models. *The Annals of Applied Statistics* 2, 1360–1383.
- Gelman, A., X.-L. Meng, and H. Stern (1996). Posterior predictive assessment of model fitness via realized discrepancies. *Statistica Sinica* 6, 733–807.
- Goldstein, M., A. Seheult, and I. Vernon (2013). Assessing model adequacy. In *Environmental Modelling: Finding Simplicity in Complexity, Second Edition* (eds J. Wainwright and M. Mulligan), pp. 435–449. John Wiley and Sons Ltd.
- Hobbs, B. P., B. P. Carlin, S. J. Mandrekar, and D. J. Sargent (2011). Hierarchical commensurate and power prior models for adaptive incorporation of historical information in clinical trials. *Biometrics* 67, 1047–1056.
- Holden, P. B., N. R. Edwards, J. Hensman, and R. D. Wilkinson (2015). ABC for climate: Dealing with expensive simulators. arXiv1511.03475.
- Ibrahim, J. G. and M.-H. Chen (2000). Power prior distributions for regression models. *Statistical Science* 15, 46–60.

- Kadane, J. B., J. M. Dickey, R. L. Winkler, W. S. Smith, and S. C. Peters (1980). Interactive elicitation of opinion for a normal linear model. *Journal of the American Statistical Association* 75, 845–854.
- Marin, J.-M., P. Pudlo, C. P. Robert, and R. Ryder (2011). Approximate Bayesian computational methods. *Statistics and Computing* 21, 289–291.
- Martin, T. G., M. A. Burgman, F. Fidler, P. M. Kuhnert, S. Low-Choy, M. McBride, and K. Mengersen (2012). Eliciting expert knowledge in conservation science. *Conservation Biology* 26(1), 29–38.
- Morris, D. E., J. E. Oakley, and J. A. Crowe (2014). A web-based tool for eliciting probability distributions from experts. *Environmental Modelling and Software* 52, 1 – 4.
- Neuenschwander, B., M. Branson, and D. J. Spiegelhalter (2009). A note on the power prior. *Statistics in Medicine* 28, 3562–3566.
- Nott, D. J., C. C. Drovandi, K. Mengersen, and M. Evans (2015). Approximation of Bayesian predictive p-values using regression ABC. *Bayesian Analysis*, to appear.
- O’Hagan, A., C. E. Buck, A. Daneshkhah, J. R. Eiser, P. H. Garthwaite, D. J. Jenkinson, J. E. Oakley, and T. Rakow (2006). *Uncertain Judgements: Eliciting Experts’ Probabilities*. Wiley.
- Pritchard, J. K., M. T. Seielstad, A. Perez-Lezaun, and M. W. Feldman (1999). Population growth of human Y chromosomes: A study of Y chromosome microsatellites. *Molecular Biology and Evolution* 16, 1791–1798.
- Racine, A., A. P. Grieve, H. Flühler, and A. F. M. Smith (1986). Bayesian methods in practice: Experiences in the pharmaceutical industry. *Journal of the Royal Statistical Society, Series C* 35, 93–150.
- Rousseeuw, P., C. Croux, V. Todorov, A. Ruckstuhl, M. Salibian-Barrera, T. Verbeke,

- M. Koller, and M. Maechler (2015). *robustbase: Basic Robust Statistics*. R package version 0.92-5.
- She, Y. and A. B. Owen (2011). Outlier detection using nonconvex penalized regression. *Journal of the American Statistical Association* 106(494).
- Silverman, B. (1986). *Density Estimation for Statistics and Data Analysis*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis.
- Simpson, D. P., H. Rue, T. G. Martins, A. Riebler, and S. H. Sørbye (2015). Penalising model component complexity: A principled, practical approach to constructing priors. arXiv1403.4630.
- Tavaré, S., D. J. Balding, R. C. Griffiths, and P. Donnelly (1997). Inferring coalescence times from DNA sequence data. *Genetics* 145(505-518).
- Venables, W. N. and B. D. Ripley (2002). *Modern Applied Statistics with S* (Fourth ed.). New York: Springer.
- Vernon, I., M. Goldstein, and R. Bower (2014). Galaxy formation: Bayesian history matching for the observable universe. *Statistical Science* 29(1), 81–90.
- Vernon, I., M. Goldstein, and R. G. Bower (2010). Galaxy formation: a Bayesian uncertainty analysis. *Bayesian Analysis* 5, 619–669.
- Wilkinson, R. (2014). Accelerating ABC methods using Gaussian processes. In *Proceedings of the Seventeenth International Conference on Artificial Intelligence and Statistics*, pp. 1015–1023.
- Williamson, D., M. Goldstein, L. Allison, A. Blaker, P. Challenor, L. Jackson, and K. Yamazaki (2013). History matching for exploring and reducing climate model parameter space using observations and a large perturbed physics ensemble. *Climate Dynamics* 41, 1703–1729.